DiCAMS White Paper

Title

DiCAMS: A Cognitive Memory Engine for Trustworthy Generative AI
By: Patrick Hoagland

Executive Abstract

DiCAMS — Dynamic Intelligent Context-Aware Memory System — is a patented, machine-learning-based cognitive memory engine that transforms stateless generative models into transparent, trustworthy, identity-consistent cognitive partners.

Today's generative AI systems, even at massive scale, lack persistent, user-governed, contextually evolving memory. They cannot preserve narrative identity, prioritize relevant knowledge, or maintain continuity over time. Retrieval-augmented generation (RAG) and scratchpad memories offer only partial workarounds, failing to support true cognitive partnership.

DiCAMS solves this missing piece. By providing a parallel "brain" that dynamically curates, organizes, prunes, and injects relevant, explainable memory into the LLM prompt stream, DiCAMS elevates generative systems from probabilistic word prediction to human-aligned, mission-critical cognitive support.

From healthcare to law, defense, education, and beyond, DiCAMS ensures generative AI can reason, remember, and adapt with continuity, privacy, and user control — the foundation of trustworthy intelligence.

---

1. Introduction

Large language models (LLMs) have transformed human-machine interaction through highly fluent text generation. But beneath their impressive performance, they remain fundamentally stateless next-token predictors. They do not maintain a personal narrative, build a lifelong memory, or reason with identity-consistent knowledge over time.

Even "memory-enabled" architectures today rely on static embeddings, ephemeral scratchpads, or large token windows — none of which provide a persistent, user-owned, dynamically evolving cognitive memory. Retrieval-augmented generation (RAG) extends

token context by injecting static chunks from a knowledge base, but cannot truly manage identity, priority, topic hierarchy, or narrative coherence.

DiCAMS directly addresses this missing link. It acts as a real-time, meta-cognitive controller of a black-box generative model:

✅ dynamic
✅ intelligent
✅ context-aware
✅ continuous
✅ human-aligned
✅ capable of correcting and steering the LLM on the fly

By adding this cognitive memory kernel, DiCAMS transforms generative models from shiny conversation toys into trusted, auditable, human-aligned cognitive partners.

---

2. Technical Background

2.1 Stateless Generative Inference

Modern LLMs operate on a sliding context window, typically measured in tokens (8k, 32k, 128k, etc.). At any given time, they see only the current window. As tokens scroll off, so does the conversation's past. This creates:

⚠️ knowledge loss
⚠️ hallucination
⚠️ narrative fragmentation
⚠️ lack of user trust

These systems cannot build a cohesive sense of "who" they are serving, nor maintain a reliable memory of long-term context, user preferences, or dynamic priorities.

2.2 Retrieval-Augmented Generation (RAG)

RAG tries to patch these gaps by retrieving semantically similar passages from a vector database, then stuffing those chunks into the prompt.

While RAG can expand context, it cannot:

⚠️ prioritize meaning
⚠️ build a narrative graph
⚠️ prune irrelevant or stale data
⚠️ preserve identity or user ownership

RAG is ultimately a search function, not a cognitive memory.

2.3 The Human Memory Analogy

Human autobiographical memory does not work like RAG or sliding windows. Humans:

✅ prioritize important memories
✅ organize them into stories
✅ maintain identity continuity
✅ prune irrelevant details
✅ update memories with new evidence

Any trustworthy, high-stakes generative AI needs to mimic these principles.

---

3. Problem Statement

LLMs today are generative engines — not memories. Their core design:

✅ token prediction
✅ maximum fluency
✅ no true knowledge retention

This leads to:

⚠️ hallucinations — inventing facts
⚠️ drift — losing track of the conversation
⚠️ lack of explainability — no traceable path
⚠️ loss of trust — especially in medicine, law, defense, and education

Without a persistent, user-governed memory structure, generative AI cannot meet the demands of high-stakes reasoning or human-centered cognitive partnership.

---

4. Proposed Solution: DiCAMS

DiCAMS — the Dynamic Intelligent Context-Aware Memory System — is a patented architecture that operates as a parallel, machine-learning-driven cognitive engine alongside any generative model.

Its functions:
✅ listen to user input in real time
✅ parse and filter relevance
✅ organize information into dynamic, hierarchical topic graphs
✅ continuously update and prune memories
✅ preserve user identity and narrative threads
✅ inject curated, trustworthy, user-controlled memory back into the LLM

The result?
DiCAMS transforms a generative model into a transparent, trustworthy, explainable, human-aligned cognitive partner.

Key Benefits Summarized:
✅ Trustworthiness — user-owned, transparent, auditable memory
✅ Explainability — traceable, replayable cognitive history
✅ Continuity — a lifelong knowledge graph
✅ Privacy — user controls what is stored, pruned, or shared
✅ Plug-and-Play — integrates with existing LLM architectures

5. DiCAMS Architecture Deep Dive

5.1 Overview

DiCAMS is a modular, machine-learning-driven cognitive memory system that operates in parallel with a generative language model. Its purpose is to:

✅ continuously parse user interactions
✅ build a user-centered knowledge graph
✅ curate, prune, and update memories dynamically

✅ inject only context-relevant, traceable, and user-approved knowledge back into the generative model

In effect, DiCAMS provides the “brain” to complement the generative “mouth.”

5.2 Core Components

Here is a clear breakdown of DiCAMS’ core architectural modules:

5.2.1 Real-Time Parser

- Actively listens to user inputs (text or voice).
- Segments information into semantic chunks.
- Labels key concepts, entities, relationships, priorities, and emotional tones.
- Filters noise, discards ephemeral chatter.

5.2.2 Topic Graph Builder

- Organizes parsed knowledge into a dynamic, hierarchical topic graph.
- Supports sub-topics, relationships, temporal sequences, and priorities.
- Mirrors human-like autobiographical memory (events, meaning, cause/effect).

5.2.3 Pruning & Conflict Resolver

- Continuously reviews the topic graph for outdated, irrelevant, or conflicting data.
- Supports user-directed corrections.
- Applies priority weighting so high-value memories persist while low-value data is pruned.

5.2.4 Identity Anchor

- Preserves a sense of user identity: preferences, roles, values, ongoing context.

- Tracks evolving self-concepts over time.
- Ensures all memory remains “about me,” coherent and user-aligned.

5.2.5 Curated Prompt Injector

- Dynamically generates an optimized prompt from the topic graph.
- Ensures the generative model sees only curated, trustworthy, user-approved context.
- Prevents token overload while maximizing context relevance.

5.2.6 Feedback Update Loop

- After each LLM response, DiCAMS captures user feedback.
- Updates the topic graph accordingly (confirming, refining, correcting, deleting nodes).
- Supports a continuously evolving cognitive framework.

5.3 End-to-End Flow

Here’s how the modules connect in real time:

1️⃣ **User speaks or types**
→ Real-Time Parser segments, tags, filters
→ Topic Graph Builder organizes context
→ Pruning module manages priority
→ Identity Anchor tracks user “who/what/when”
→ Curated Prompt Injector sends high-fidelity, optimized knowledge to the LLM
→ LLM generates a response
→ Feedback Update Loop corrects, reinforces, or expands the topic graph
→ repeat

5.4 Engineering Principles

✅ *Parallel cognitive architecture* — DiCAMS does not replace the LLM, it co-evolves beside it.

✅ *User-governed provenance* — every memory is transparent, timestamped, user-inspectable.

✅ *Real-time performance* — built for live, interactive, human-paced conversations.

✅ *Modular extensibility* — pluggable with any LLM, regardless of vendor or framework.

✅ *Privacy-first* — no data leaves the user's ownership unless explicitly approved.

5.5 Visual Diagram (Textual Sketch)

```pgsql
USER INPUT
   ↓
[Real-Time Parser]
   ↓
[Topic Graph Builder]
   ↓
[Pruning / Conflict Resolver]
   ↓
[Identity Anchor]
   ↓
[Curated Prompt Injector]
   ↓
      [Generative LLM]
   ↓
[Feedback Update Loop]
   ↺ (back to Topic Graph)
```

👉 *Think of it as a feedback spiral — each cycle grows more consistent, relevant, and user-aligned.*

6. Pilot Use Cases

DiCAMS is designed as a universal cognitive memory engine. However, its value is clearest in domains where:

✅ trust is critical
✅ errors can be catastrophic
✅ memory continuity is essential

Here is a polished breakdown of its priority pilot targets:

6.1 Healthcare

✅ **Problem:** Doctors and patients exchange high-volume, high-stakes information. Current EHR systems and clinical notetaking are static and error-prone. Hallucinations in generative models present unacceptable risks.

✅ **DiCAMS Impact:**

- Real-time conversation capture, structured into evolving patient histories.
- Dynamic topic graphs of symptoms, labs, meds, and prior diagnoses.
- Curated, timestamped, traceable prompts to the LLM, minimizing hallucination.
- Full transparency for audit, clinical review, and shared decision-making.

✅ **Result:** A true cognitive partner that preserves patient context over multiple visits, improving care continuity, explainability, and trust.

6.2 Legal

✅ **Problem:** Lawyers handle thousands of documents, shifting arguments, complex precedents, and changing case strategies. LLMs risk hallucinating citations or losing context.

✅ **DiCAMS Impact:**

- Structured, dynamic knowledge graph of case details, filings, evidence, and witness statements.
- Maintains consistent topic hierarchies across weeks or years.

- Prunes irrelevant or outdated arguments, maintaining clarity.
- Feeds curated, traceable prompts to the LLM for legal co-piloting.

✅ **Result:** Higher legal confidence, reduced errors, and an explainable memory of every argument and ruling.

6.3 Defense

✅ **Problem:** Mission-critical information in defense and security cannot risk hallucinations or context drift. Chain-of-command consistency is paramount.

✅ **DiCAMS Impact:**

- Captures orders, situational updates, mission parameters, and security clearances in a structured graph.
- Maintains a dynamic, consistent mission thread over time.
- Provides transparent, traceable memory inputs to generative co-pilots.

✅ **Result:** Trustworthy human-AI collaboration with a transparent audit trail — essential for defense-grade decision support.

6.4 Education

✅ **Problem:** Students and educators struggle with context loss over semesters, fragmented notes, and knowledge gaps.

✅ **DiCAMS Impact:**

- Builds a dynamic, evolving graph of each learner's knowledge journey.
- Supports persistent context across lessons, semesters, or years.
- Allows transparent review of mistakes, growth, and areas to reinforce.
- Feeds curated, trustworthy context to tutoring co-pilots.

✅ **Result:** A personalized cognitive tutor that evolves with the student, making learning explainable, persistent, and growth-oriented.

6.5 Universal Applications

✅ **Problem:** LLMs in general-purpose applications still suffer from:

- context window limits
- knowledge drift
- hallucination
- lack of user ownership

✅ **DiCAMS Impact Across Domains:**

- X-rays or radiology: persistent case memory
- Financial audits: consistent anomaly tracking
- Historical archives: trustworthy timeline reconstruction
- Group knowledge management: dynamic, consensus-building memory
- Personal knowledge: lifelong, user-owned cognitive thread

✅ **Result:** Any user — from an individual knowledge worker to a global team — benefits from a trusted, explainable, contextually consistent generative experience.

7. Competitive Landscape

7.1 Today's Memory Approaches in Generative AI

Most current generative systems attempt to solve memory limitations with one or more of these approaches:

✅ **Retrieval-Augmented Generation (RAG)**

- Uses a vector store to fetch semantically similar documents
- Injects static chunks into the prompt
- Fundamentally keyword-based, with no evolving narrative graph
- Offers no true prioritization, no topic pruning, no user-governed structure

✅ **Scratchpad Memory**

- Keeps ephemeral chat transcripts or token archives

- Provides no identity-consistent narrative
- Cannot organize knowledge with hierarchy or priority
- Becomes bloated and unmanageable over time

✅ **Vector-Based Memory**

- Embeds user facts for retrieval
- Ignores time, causality, topic evolution
- No user traceability or explainable pruning

✅ **Agent Orchestrators**

- Chain multiple tools and agents to create complex workflows
- But still fundamentally rely on stateless token histories
- No centralized, evolving autobiographical memory

7.2 The DiCAMS Difference

DiCAMS is fundamentally different. It does not "retrieve" or "snapshot"; it **co-evolves** with the user:

✅ **Dynamic** — parses, updates, and prunes memory live, rather than passively storing it
✅ **Hierarchical** — builds a topic graph with relationships, priorities, and time-sequence
✅ **User-Aligned** — memories are user-owned, transparent, and auditable
✅ **Identity-Preserving** — maintains personal preferences, roles, and values across time
✅ **Explainable** — every memory injection is traceable, with a clear provenance
✅ **Plug-and-Play** — modular, compatible with any generative engine

No other current architecture delivers:

🟢 a living, curated, narrative memory graph
🟢 a user-controlled cognitive prosthetic
🟢 built-in priority pruning
🟢 timestamped, transparent topic continuity
🟢 identity-consistent interaction

DiCAMS transforms the generative black box into a collaborative, context-aware, trustworthy cognitive partner — rather than a static or ephemeral assistant.

7.3 Strategic Advantage

✅ **Patented** — DiCAMS is already protected under granted U.S. intellectual property, making it defensible against copycats.

✅ **Modular** — can integrate with any LLM architecture, from OpenAI to open-source models, providing rapid deployment options.

✅ **Future-Proof** — as LLM token windows expand, DiCAMS scales its curated injection approach, avoiding the bloated "sticky note stadium" problem and keeping the knowledge clean, relevant, and explainable.

✅ **First-to-Market Opportunity** — while other players focus on retrieval patches, DiCAMS provides a fundamentally novel, human-inspired cognitive framework that investors and enterprise partners can rally behind.

8. Benefits of DiCAMS

8.1 Trustworthiness

✅ By providing a transparent, user-owned, timestamped cognitive memory, DiCAMS dramatically increases the trustworthiness of generative systems.
✅ Users and auditors can see exactly what knowledge was used, how it evolved, and where it came from — eliminating black-box hallucinations and opaque reasoning.

8.2 Explainability

✅ DiCAMS makes every LLM output explainable by maintaining a structured, traceable memory graph.
✅ Any generated answer can be mapped back to the curated topic graph, revealing precisely how the LLM was guided to its response.

8.3 Continuity

✅ Unlike ephemeral or snapshot-based memory, DiCAMS maintains a persistent, user-aligned, evolving cognitive thread across conversations, days, or even years.
✅ This continuity mirrors human autobiographical memory, providing the LLM with stable, consistent context to reason more like a trusted human partner.

8.4 Privacy and Governance

✅ DiCAMS is built around a privacy-first, user-governed model.
✅ Users decide what to store, what to delete, and what to share, maintaining full control of their cognitive footprint.
✅ This meets — and exceeds — GDPR, HIPAA, and other high-trust compliance frameworks.

8.5 Compatibility

✅ DiCAMS is a modular layer, designed to integrate with any LLM pipeline regardless of vendor or hosting environment.
✅ This allows rapid adoption in existing architectures without expensive rebuilds.

8.6 Cognitive Offloading

✅ DiCAMS acts as a trusted cognitive prosthetic, offloading complex memory, prioritization, summarization, and error correction from the human user.
✅ Professionals — doctors, lawyers, researchers — can focus on high-value reasoning, while DiCAMS ensures consistent, reliable, explainable recall.

8.7 Human-Centric Collaboration

✅ DiCAMS enables a new paradigm where generative systems behave more like an *augmented cognition partner* than a passive chatbot.
✅ The LLM becomes a co-reasoning companion, informed by the user's values, narrative, and evolving mental model.

8.8 Market Advantage

✅ Combining trust, explainability, and user-governed memory positions DiCAMS as the first true cognitive memory engine for generative AI.
✅ This unlocks markets (healthcare, defense, legal, education) that otherwise cannot adopt generative systems because of trust and compliance concerns.

9. Implementation Roadmap

9.1 Stage 1: Visual Simulation (0–3 Months)

✅ **Goal:** Communicate the DiCAMS vision credibly and compellingly to investors, partners, and pilot customers.

- Develop a high-fidelity explainer video showing a realistic, scenario-driven animation (e.g., a doctor–patient conversation) with DiCAMS operating in real time.
- Simulate how memory is chunked, structured, prioritized, and re-injected into the LLM.
- Include visual overlays showing user-governed knowledge graphs and memory pruning.
- Demonstrate the "why" behind each memory chunk being fed to the LLM.

✅ *Outcome:* An investor-grade visual artifact demonstrating exactly how DiCAMS works — even before live code.

9.2 Stage 2: Minimal Viable Pilot (3–9 Months)

✅ **Goal:** Build a functional demonstration of DiCAMS in a tightly scoped pilot.

- Focus on a high-trust domain: *healthcare* or *legal*.
- Use limited data sources, e.g., synthetic or de-identified patient records or case files.
- Implement real-time parsing, graph construction, memory pruning, and prompt injection.
- Prove a closed memory loop with user-governed context and explainable output.
- Establish auditing and traceability workflows.

✅ *Outcome:* A working prototype showing that DiCAMS reliably creates explainable, user-owned cognitive memory in a live setting.

9.3 Stage 3: Commercial Licensing or Scale (9–18 Months)

✅ **Goal:** Move beyond pilot to production-grade integration.

- License the DiCAMS kernel to foundation model providers or vertical AI platforms (e.g., EHR systems, legal tech, defense systems).
- Integrate DiCAMS as a modular cognitive memory backbone with commercial-grade security, compliance, and performance.
- Build a revenue-sharing or licensing model with strategic partners.
- Expand pilot learnings into new verticals (education, financial compliance, personal knowledge management).

✅ *Outcome:* DiCAMS becomes a defensible, revenue-generating platform licensed to enterprise-grade generative AI systems.

9.4 Parallel Research Track

✅ In parallel, DiCAMS can initiate a research program to:

- Collaborate with universities and research hospitals to validate trauma-informed AI and narrative-based memory in mental health or PTSD contexts.
- Publish findings on human–AI cognitive partnership.
- Build an academic community around DiCAMS as a new field of cognitive AI augmentation.

✅ *Outcome:* Scientific credibility, peer-reviewed research, and a potential path to grant funding.

9.5 Governance and Ethics

✅ Build governance frameworks early, including:

- User data privacy controls
- Memory retention/expiration policies
- Consent-driven memory updates
- GDPR and HIPAA compliance models

✅ *Outcome:* Position DiCAMS as *the* most trustworthy and ethical memory framework for generative AI, protecting user rights from day one.

10. Business Model & Licensing

10.1 Licensing Strategy

✅ **Kernel Licensing**

- DiCAMS will be offered as a modular, patent-protected "kernel" that can integrate with existing LLM infrastructures.
- Foundation model vendors (e.g., OpenAI, Anthropic, Cohere) or vertical AI platforms (e.g., Epic EHR, defense contractors) can license the DiCAMS kernel to embed trustworthy memory capabilities directly into their stack.

✅ **Revenue Model**

- Per-seat licensing for enterprise installations
- Per-query transaction fees for high-volume cloud APIs
- Platform royalties for third-party integrations
- Joint-venture co-development for specialized verticals (e.g., military, government, education)

✅ *Outcome:* Predictable, recurring revenue streams, high-margin software, and defensible IP-based growth.

10.2 Professional Services

✅ **Deployment Support**

- Technical integration assistance

- Memory policy design workshops (e.g., GDPR/HIPAA compliance)
- Training programs for user-governed memory best practices

✅ **Consulting**

- Custom architecture designs for unique high-trust use cases
- Cognitive memory audits and risk assessments

✅ *Outcome:* Professional services build stickiness with customers, create long-term relationships, and expand DiCAMS' reach as the de facto memory standard.

10.3 Research Collaborations

✅ DiCAMS can partner with universities, clinical research labs, or government agencies to:

- Validate the cognitive memory architecture through published trials
- Secure grant funding
- Create a public–private innovation pipeline

✅ *Outcome:* Enhances credibility, demonstrates social benefit, and attracts mission-driven partners.

10.4 Strategic Partnerships

✅ **Enterprise SaaS & Cloud**

- Collaborate with hyperscale cloud vendors to provide DiCAMS as an enterprise-ready plugin
- Use DiCAMS to differentiate their generative offerings in regulated markets

✅ **Domain-Specific Integrators**

- Partner with EHR vendors, legal case management systems, defense intelligence platforms
- Integrate DiCAMS as a cognitive co-pilot framework that is trustable and compliant

✅ *Outcome:* Expands market footprint rapidly, leveraging trusted enterprise distribution.

10.5 Competitive Moat

✅ **Patent Protection**

- DiCAMS is already granted U.S. patent protection, defending its novel cognitive memory architecture from replication.

✅ **First-Mover Advantage**

- There is no equivalent plug-and-play cognitive memory engine on the market today.
- Early partnerships and licensing deals will establish DiCAMS as the gold standard.

✅ **Brand Trust**

- Position DiCAMS as *the* solution for human-centered, transparent, explainable memory in generative AI.
- Emphasize user-owned memory and cognitive agency.

✅ *Outcome:* A robust, defensible competitive moat, locking in strategic market share.

11. Risks & Mitigation

11.1 Technical Risk

⚠️ **Risk**

- Complexity of integrating DiCAMS with diverse LLM architectures
- Potential performance bottlenecks in real-time memory graph parsing and injection

✅ **Mitigation**

- Design a modular API with well-documented interfaces
- Start with a narrowly scoped healthcare or legal pilot before general rollout
- Build caching and asynchronous memory update pipelines to manage latency

---

11.2 Adoption Risk

⚠️ **Risk**

- Resistance from large foundation model vendors to incorporate external cognitive layers
- Potential cultural bias in trusting “memory augmentation”

✅ **Mitigation**

- Provide plug-and-play SDKs with minimal architectural disruption
- Co-market with vertical partners who have the most to gain (e.g., hospitals, law firms)
- Showcase pilot case studies that demonstrate trustworthiness and real-world ROI

11.3 Privacy & Data Risk

⚠️ **Risk**

- Sensitive user memories require robust data security, retention, and deletion policies
- Regulatory hurdles (GDPR, HIPAA, defense clearances)

✅ **Mitigation**

- Build privacy-by-design features from day one
- Incorporate robust consent management and user data ownership controls
- Work with external privacy auditors and legal counsel to align with compliance frameworks

---

11.4 Competitive Risk

⚠️ **Risk**

- Other vendors may attempt to copy DiCAMS’ features
- Foundation model providers might try to build internal memory solutions

✅ **Mitigation**

- Assert and actively defend DiCAMS’ granted patent
- Continue to innovate faster with superior user-centered design

- Establish first-to-market licensing deals and partnerships to secure territory early

---

11.5 Ethical Risk

⚠️ **Risk**

- Misuse of memory features (e.g., biased memory curation, user manipulation)
- Concerns about automating too much of personal cognition

✅ **Mitigation**

- Enforce explainable, user-owned memory frameworks
- Provide robust auditing, transparency logs, and opt-out options
- Engage with ethics boards and independent oversight to preserve trust

---

11.6 Operational Risk

⚠️ **Risk**

- Startup resource constraints
- Talent recruitment challenges
- Coordination across partners

✅ **Mitigation**

- Prioritize an MVP with a small, high-quality team
- Lean on strategic advisors (e.g., Marc) for introductions to partners
- Structure staged financing milestones tied to pilot outcomes

12. Conclusion & Next Steps

12.1 Conclusion

DiCAMS — the Dynamic Intelligent Context-Aware Memory System — represents a fundamental leap beyond today's generative AI. Where existing systems operate with ephemeral, stateless, and opaque memory, DiCAMS introduces a transparent, user-governed, explainable cognitive memory engine.

By acting as a real-time, meta-cognitive controller of black-box language models, DiCAMS:

✅ preserves personal meaning and narrative identity

✅ offloads cognitive burden

✅ ensures trustworthy, auditable knowledge flow

✅ aligns generative AI with human-centered values

It transforms any LLM into a genuine cognitive partner — consistent, transparent, and safe enough for high-stakes industries. Healthcare, law, defense, education, and any trust-critical domain can now benefit from generative AI with confidence, thanks to DiCAMS.

12.2 Next Steps

👉 Visual Simulation — Begin development of a high-fidelity explainer video showing DiCAMS in action within a realistic scenario (e.g., doctor–patient conversation).

👉 Pilot Deployment — Secure a pilot partner in a trust-heavy domain (healthcare or legal), to validate DiCAMS with real-world data and user feedback.

👉 Investor Outreach — Share this white paper, the visual simulation, and pilot roadmap with selected investors, demonstrating clear product-market fit and defensible IP.

👉 Strategic Evangelist — Identify and engage a technically fluent, mission-aligned evangelist who can present DiCAMS to the world with clarity and passion.

👉 Community & Research — Initiate partnerships with academic and clinical researchers to validate DiCAMS' benefits, build credibility, and contribute to a growing field of human-aligned cognitive augmentation.

12.3 Final Note

DiCAMS is more than a technical invention — it is a breakthrough in human–AI collaboration, built to restore agency, trust, and narrative coherence in an age of black-box generative models.

✅ The time to bring it to market is now.

✅ The world is demanding explainable, safe, trustworthy AI.

✅ DiCAMS is ready.